

# Exhibit 1

# Exhibit 2

**U.S. Patent No. 7,519,814 vs. Microsoft**

Accused Instrumentalities: Microsoft products and services using secure containerized applications, including without limitation Azure Kubernetes Service (“AKS”), Azure Arc-enabled Kubernetes, Azure Container Registry, and Azure Container Apps, and all versions and variations thereof since the issuance of the asserted patent.

**Claim 1**

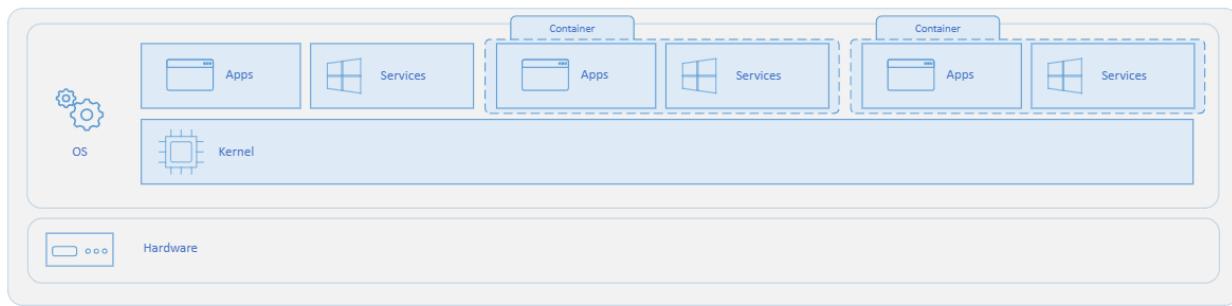
Claim 1	Accused Instrumentalities
<p>[1pre] 1. In a system having a plurality of servers with operating systems that differ, operating in disparate computing environments, wherein each server includes a processor and an operating system including a kernel a set of associated local system files compatible with the processor, a method of providing at least some of the servers in the system with secure, executable, applications related to a service, wherein the applications are executed in a secure environment, wherein the applications each include an object executable by at least some of the different operating systems for performing a task related to the service, the method comprising:</p>	<p>To the extent the preamble is limiting, Microsoft and/or its customer practices, through the Accused Instrumentalities, in a system having a plurality of servers with operating systems that differ, operating in disparate computing environments, wherein each server includes a processor and an operating system including a kernel a set of associated local system files compatible with the processor, a method of providing at least some of the servers in the system with secure, executable, applications related to a service, wherein the applications are executed in a secure environment, wherein the applications each include an object executable by at least some of the different operating systems for performing a task related to the service, as claimed.</p> <p>For example, Azure Kubernetes Service runs on individual servers, each of which runs an independent operating system running either on bare metal, through an on-premises virtualized infrastructure, through one or more cloud services, or through any other supported deployment. In an exemplary deployment, two or more servers use different operating systems. The servers operate in disparate computing environments, including because each server is a stand-alone computer and/or each server is unrelated to the other servers due to having independent hardware and, in some instances, independent software.</p> <p>Microsoft requires and/or provides that each server includes a processor with one or more cores available to the OS kernel. Microsoft further requires and/or provides that each server has a supported operating system, which includes a kernel and associated local system files, configuration files, etc. In the infringing system, at least two servers have different operating systems.</p> <p>In at least some instances, Microsoft directly owns, operates, controls, and/or benefits from the claimed system and/or method. In other instances, Microsoft’s customer makes and uses the system and/or method either by following Microsoft’s direction and control, including Microsoft’s documentation, or automatically through the ordinary and expected operation of Microsoft’s software, or a combination thereof.</p>

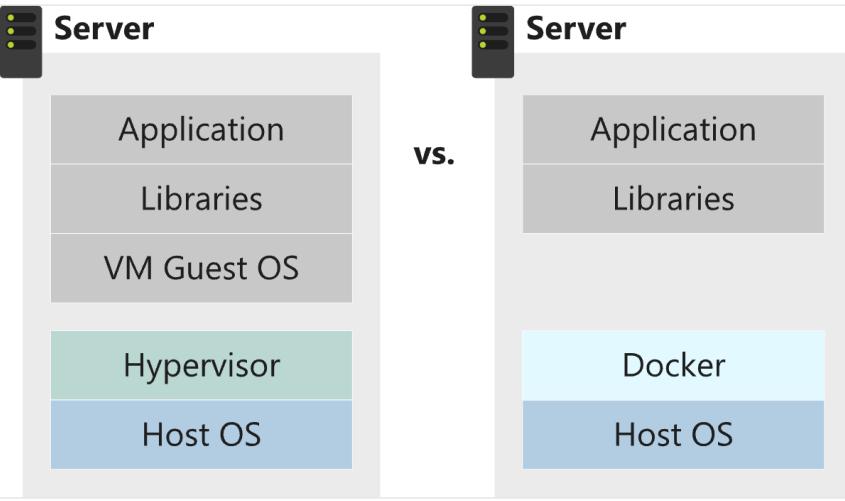
Claim 1	Accused Instrumentalities
	<p><i>See</i> claim limitations below.</p> <p><i>See also</i>, e.g.:</p> <p>Azure Kubernetes Service (AKS) is a managed Kubernetes service that you can use to deploy and manage containerized applications. You need minimal container orchestration expertise to use AKS. AKS reduces the complexity and operational overhead of managing Kubernetes by offloading much of that responsibility to Azure. AKS is an ideal platform for deploying and managing containerized applications that require high availability, scalability, and portability, and for deploying applications to multiple regions, using open-source tools, and integrating with existing DevOps tools.</p> <p><a href="https://learn.microsoft.com/en-us/azure/aks/what-is-aks">https://learn.microsoft.com/en-us/azure/aks/what-is-aks</a></p> <h2>When to use AKS</h2> <p>The following list describes some common use cases for AKS:</p> <ul style="list-style-type: none"> <li>• <b>Lift and shift to containers with AKS:</b> Migrate existing applications to containers and run them in a fully managed Kubernetes environment.</li> <li>• <b>Microservices with AKS:</b> Simplify the deployment and management of microservices-based applications with streamlined horizontal scaling, self-healing, load balancing, and secret management.</li> <li>• <b>Secure DevOps for AKS:</b> Efficiently balance speed and security by implementing secure DevOps with Kubernetes.</li> <li>• <b>Bursting from AKS with ACI:</b> Use virtual nodes to provision pods inside ACI that start in seconds and scale to meet demand.</li> <li>• <b>Machine learning model training with AKS:</b> Train models using large datasets with familiar tools, such as TensorFlow and Kubeflow.</li> <li>• <b>Data streaming with AKS:</b> Ingest and process real-time data streams with millions of data points collected via sensors, and perform fast analyses and computations to develop insights into complex scenarios.</li> <li>• <b>Using Windows containers on AKS:</b> Run Windows Server containers on AKS to modernize your Windows applications and infrastructure.</li> </ul> <p><a href="https://learn.microsoft.com/en-us/azure/aks/what-is-aks">https://learn.microsoft.com/en-us/azure/aks/what-is-aks</a></p>

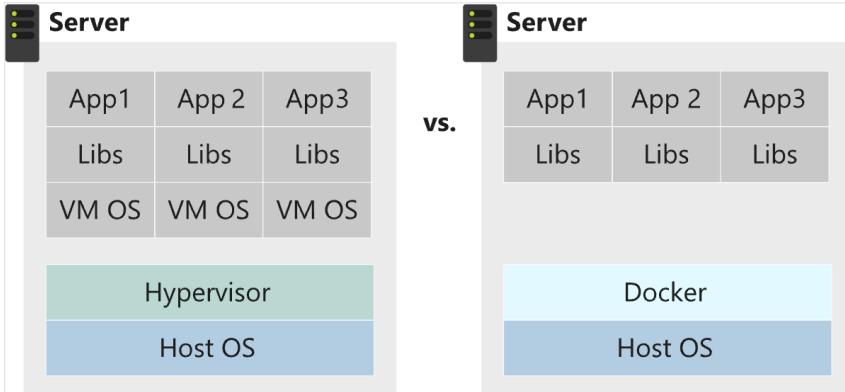
Claim 1	Accused Instrumentalities
	<p>Azure Arc-enabled Kubernetes allows you to attach Kubernetes clusters running anywhere so that you can manage and configure them in Azure. By managing all of your Kubernetes resources in a single control plane, you can enable a more consistent development and operation experience, helping you run cloud-native apps anywhere and on any Kubernetes platform.</p> <p>When the <a href="#">Azure Arc agents are deployed to the cluster</a>, an outbound connection to Azure is initiated, using industry-standard SSL to secure data in transit.</p> <p>Clusters that you connect to Azure are represented as their own resources in Azure Resource Manager, and they can be organized using resource groups and tagging.</p> <p><a href="https://learn.microsoft.com/en-us/azure/azure-arc/kubernetes/overview">https://learn.microsoft.com/en-us/azure/azure-arc/kubernetes/overview</a></p> <p>Containers are becoming the preferred way to package, deploy, and manage cloud applications. Azure Container Instances offers the fastest and simplest way to run Linux or Windows containers in Azure, without having to manage any virtual machines and without having to adopt a higher-level service.</p> <p>ACI supports <a href="#">regular</a>, <a href="#">confidential</a>, and <a href="#">Spot</a> containers. ACI can be used as <a href="#">single-instance</a> or multi-instance via <a href="#">NGroups</a>, or you can get orchestration capabilities by deploying pods in your Azure Kubernetes Service (AKS) cluster via <a href="#">virtual nodes on ACI</a>. For even faster startup times, ACI supports <a href="#">standby pools</a>.</p> <p><a href="https://learn.microsoft.com/en-us/azure/container-instances/container-instances-overview">https://learn.microsoft.com/en-us/azure/container-instances/container-instances-overview</a></p>

Claim 1	Accused Instrumentalities
	<p>Azure Container Apps is a serverless platform that allows you to maintain less infrastructure and save costs while running containerized applications. Instead of worrying about server configuration, container orchestration, and deployment details, Container Apps provides all the up-to-date server resources required to keep your applications stable and secure.</p> <p>Common uses of Azure Container Apps include:</p> <ul style="list-style-type: none"> <li>• Deploying API endpoints</li> <li>• Hosting background processing jobs</li> <li>• Handling event-driven processing</li> <li>• Running microservices</li> </ul> <p><a href="https://learn.microsoft.com/en-us/azure/container-apps/overview">https://learn.microsoft.com/en-us/azure/container-apps/overview</a></p>

Claim 1	Accused Instrumentalities			
	<p>Applies to: Windows Server 2022, Windows Server 2019, Windows Server 2016</p> <p>Containers are a technology for packaging and running Windows and Linux applications across diverse environments on-premises and in the cloud. Containers provide a lightweight, isolated environment that makes apps easier to develop, deploy, and manage. Containers start and stop quickly, making them ideal for apps that need to rapidly adapt to changing demand. The lightweight nature of containers also make them a useful tool for increasing the density and utilization of your infrastructure.</p> <div data-bbox="699 535 1579 1073"> <table border="0"> <tbody> <tr> <td data-bbox="699 535 988 1073"> <p><b>Anywhere</b></p>  <p>On-premises</p>  <p>Cloud</p> </td> <td data-bbox="988 535 1277 1073"> <p><b>Any app</b></p>  <p>Monolith</p>  <p>Microservice</p> </td> <td data-bbox="1277 535 1579 1073"> <p><b>Any language</b></p>  <p>Java</p>  <p>.Net</p>  <p>Python</p>  <p>Node</p> </td> </tr> </tbody> </table> </div> <p><a href="https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/">https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/</a></p>	<p><b>Anywhere</b></p>  <p>On-premises</p>  <p>Cloud</p>	<p><b>Any app</b></p>  <p>Monolith</p>  <p>Microservice</p>	<p><b>Any language</b></p>  <p>Java</p>  <p>.Net</p>  <p>Python</p>  <p>Node</p>
<p><b>Anywhere</b></p>  <p>On-premises</p>  <p>Cloud</p>	<p><b>Any app</b></p>  <p>Monolith</p>  <p>Microservice</p>	<p><b>Any language</b></p>  <p>Java</p>  <p>.Net</p>  <p>Python</p>  <p>Node</p>		

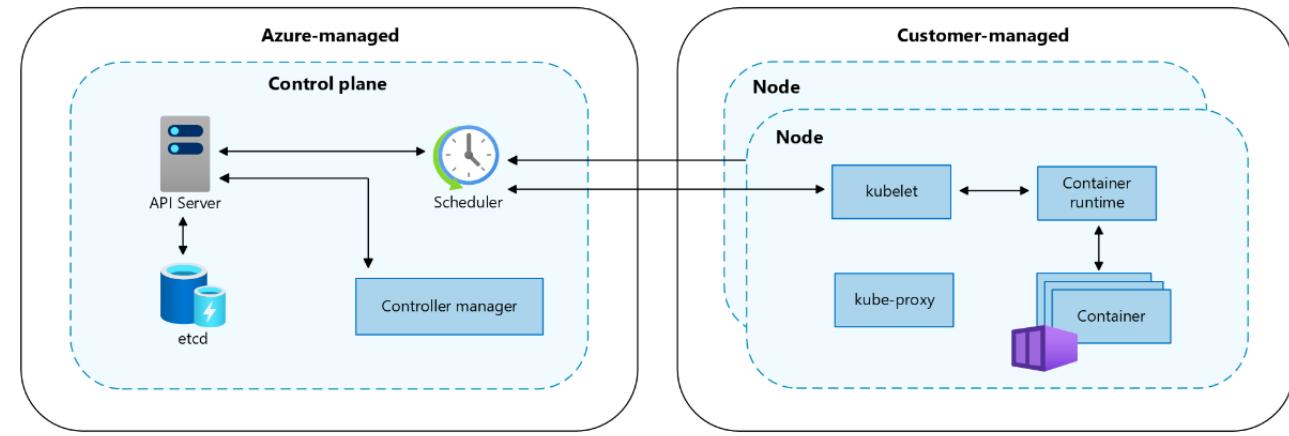
Claim 1	Accused Instrumentalities
	<h2 data-bbox="692 230 1227 279">How containers work</h2> <p data-bbox="692 319 1981 434">A container is an isolated, lightweight package for running an application on the host operating system. Containers build on top of the host operating system's kernel (which can be thought of as the buried plumbing of the operating system), as shown in the diagram below.</p>  <p data-bbox="692 842 1981 1037">While a container shares the host operating system's kernel, the container doesn't get unfettered access to it. Instead, the container gets an isolated—and in some cases virtualized—view of the system. For example, a container can access a virtualized version of the file system and registry, but any changes affect only the container and are discarded when it stops. To save data, the container can mount persistent storage such as an <a href="#">Azure Disk</a> or a file share (including <a href="#">Azure Files</a>).</p> <p data-bbox="692 1070 1981 1307">A container builds on top of the kernel, but the kernel doesn't provide all of the APIs and services an app needs to run—most of these are provided by system files (libraries) that run above the kernel in user mode. Because a container is isolated from the host's user mode environment, the container needs its own copy of these user mode system files, which are packaged into something known as a base image. The base image serves as the foundational layer upon which your container is built, providing it with operating system services not provided by the kernel. But we'll talk more about container images later.</p> <p data-bbox="675 1323 1628 1364"><a href="https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/">https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/</a></p>

Claim 1	Accused Instrumentalities
	<p><b>Docker benefits</b></p> <p>When we use Docker, we immediately get access to the benefits containerization offer.</p> <p><b>Efficient hardware use</b></p> <p>Containers run without using a virtual machine (VM). As we learned, the container relies on the host kernel for functions such as file system, network management, process scheduling, and memory management.</p>  <p>Compared to a VM, we can see that a VM requires an OS installed to provide kernel functions to the running applications inside the VM. Keep in mind that the VM OS also requires disk space, memory, and CPU time. By removing the VM and the additional OS requirement, we can free resources on the host and use it for running other containers.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p>

Claim 1	Accused Instrumentalities
	<p><b>Container isolation</b></p> <p>Docker containers provide security features to run multiple containers simultaneously on the same host without affecting each other. As we learned, we can configure both data storage and network configuration to isolate our containers or share data and connectivity between specific containers.</p> <p>Let's compare this feature to using VMs.</p>  <p>Assume we have a physical host running two VMs. We have three applications that we want to run isolated from each other. We decide to deploy the first app onto VM1 and the second onto VM2 to separate the two apps from each other. If we now choose to install the third application, we'll need to install another VM to continue this pattern.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p> <p><b>Application portability</b></p> <p>Containers run almost everywhere: desktops, physical servers, VMs, and in the cloud. This runtime compatibility makes it easy to move containerized applications among different environments.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p>

Claim 1	Accused Instrumentalities
	<p><b>Cloud deployments</b></p> <p>Docker containers are the default container architecture the Azure containerization services use, and many other cloud platforms also support them.</p> <p>For instance, you can deploy Docker containers to Azure Container Instances, Azure App Service, and Azure Kubernetes Services. Each of these options provides you with different features and capabilities.</p> <p>For example, Azure container instances allow you to focus on designing and building your applications without the overhead of managing infrastructure. When you have many containers to orchestrate, Azure Kubernetes service makes it easy to deploy and manage large-scale container deployments.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p> <p>There are many different orchestrators that you can use with Windows containers; here are the options Microsoft provides:</p> <ul style="list-style-type: none"> <li>• <a href="#">Azure Kubernetes Service (AKS)</a> - use a managed Azure Kubernetes service</li> <li>• <a href="#">Azure Kubernetes Service (AKS) on Azure Stack HCI</a> - use Azure Kubernetes Service on-premises</li> </ul> <p><a href="https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/">https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/</a></p>

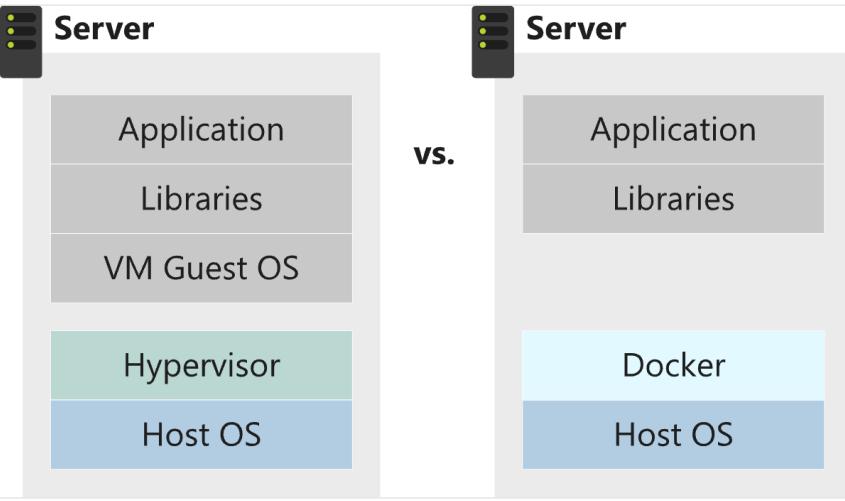
Claim 1	Accused Instrumentalities
	<p><b>What is Kubernetes?</b></p> <p>Kubernetes is an open-source container orchestration platform for automating the deployment, scaling, and management of containerized applications. For more information, see the official <a href="#">Kubernetes documentation</a>.</p> <p><b>What is AKS?</b></p> <p>AKS is a managed Kubernetes service that simplifies deploying, managing, and scaling containerized applications using Kubernetes. For more information, see <a href="#">What is Azure Kubernetes Service (AKS)?</a></p> <p><a href="https://learn.microsoft.com/en-us/azure/aks/core-aks-concepts">https://learn.microsoft.com/en-us/azure/aks/core-aks-concepts</a></p>

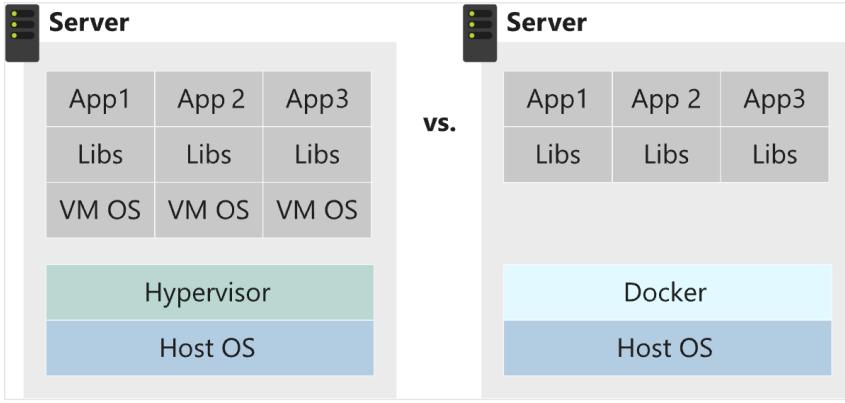
Claim 1	Accused Instrumentalities
	<h2 data-bbox="677 208 1184 267">Cluster components</h2> <p data-bbox="677 300 1290 328">An AKS cluster is divided into two main components:</p> <ul data-bbox="720 368 1930 479" style="list-style-type: none"> <li data-bbox="720 368 1930 430">• <b>Control plane:</b> The control plane provides the core Kubernetes services and orchestration of application workloads.</li> <li data-bbox="720 447 1691 479">• <b>Nodes:</b> Nodes are the underlying virtual machines (VMs) that run your applications.</li> </ul>  <p>The diagram illustrates the architecture of an AKS cluster. It is divided into two main sections: the <b>Azure-managed Control plane</b> and the <b>Customer-managed Node</b>.</p> <p><b>Azure-managed Control plane:</b> This section is enclosed in a dashed blue box. It contains the <b>API Server</b> (represented by a grey server icon), the <b>Scheduler</b> (represented by a clock icon), and the <b>Controller manager</b> (represented by a blue rectangle). The API Server is connected to the Scheduler and the Controller manager. The Scheduler is connected to the Controller manager. The Controller manager is connected to the API Server.</p> <p><b>Customer-managed Node:</b> This section is also enclosed in a dashed blue box. It contains the <b>kubelet</b> (represented by a blue rectangle), the <b>Container runtime</b> (represented by a blue rectangle), the <b>kube-proxy</b> (represented by a blue rectangle), and the <b>Container</b> (represented by a purple cube icon). The kubelet is connected to the Container runtime. The Container runtime is connected to the Container. The kube-proxy is connected to the Container.</p> <p><a href="https://learn.microsoft.com/en-us/azure/aks/core-aks-concepts">https://learn.microsoft.com/en-us/azure/aks/core-aks-concepts</a></p>

Claim 1	Accused Instrumentalities
	<p><b>VM size and image</b></p> <p>The Azure VM size for your nodes defines CPUs, memory, size, and the storage type available, such as high-performance SSD or regular HDD. The VM size you choose depends on the workload requirements and the number of pods you plan to run on each node. For more information, see <a href="#">Supported VM sizes in Azure Kubernetes Service (AKS)</a>.</p> <p>In AKS, the VM image for your cluster's nodes is based on Ubuntu Linux, <a href="#">Azure Linux</a>, or Windows Server 2022. When you create an AKS cluster or scale out the number of nodes, the Azure platform automatically creates and configures the requested number of VMs. Agent nodes are billed as standard VMs, so any VM size discounts, including <a href="#">Azure reservations</a>, are automatically applied.</p> <p><a href="https://learn.microsoft.com/en-us/azure/aks/core-aks-concepts">https://learn.microsoft.com/en-us/azure/aks/core-aks-concepts</a></p> <p>In order to run Windows containers, your Kubernetes cluster must include multiple operating systems. While you can only run the <a href="#">control plane</a> on Linux, you can deploy worker nodes running either Windows or Linux.</p> <p>Windows <a href="#">nodes</a> are <a href="#">supported</a> provided that the operating system is Windows Server 2019 or Windows Server 2022.</p> <p>This document uses the term <i>Windows containers</i> to mean Windows containers with process isolation. Kubernetes does not support running Windows containers with <a href="#">Hyper-V isolation</a>.</p> <p><a href="https://kubernetes.io/docs/concepts/windows/intro/">https://kubernetes.io/docs/concepts/windows/intro/</a></p>

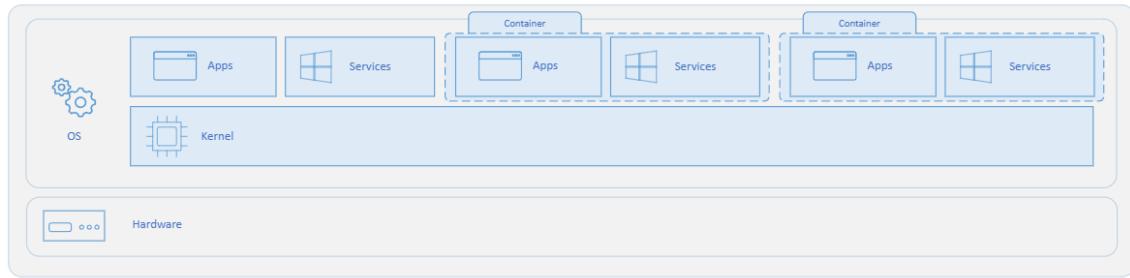
Claim 1	Accused Instrumentalities
	<p><b>Are there any limitations on the number of services on a cluster with Windows nodes?</b></p> <p>A cluster with Windows nodes can have approximately 500 services (sometimes less) before it encounters port exhaustion. This limitation applies to a Kubernetes Service with External Traffic Policy set to "Cluster".</p> <p>When the external traffic policy on a Service is configured as a Cluster, the traffic undergoes an extra Source NAT on the node. This process also results in reservation of a port from the TCPIP dynamic port pool. This port pool is a limited resource (~16K ports by default) and many active connections to a Service can lead to dynamic port pool exhaustion resulting in connection drops.</p> <p>If the Kubernetes Service is configured with External Traffic Policy set to "Local", port exhaustion problems aren't likely to occur at 500 services.</p> <p><a href="https://learn.microsoft.com/en-us/azure/aks/windows-faq">https://learn.microsoft.com/en-us/azure/aks/windows-faq</a></p> <p>Kernel mode refers to the processor mode that enables software to have full and unrestricted access to the system and its resources. The OS kernel and kernel drivers, such as the file system driver, are loaded into protected memory space and operate in this highly privileged kernel mode.</p> <p><a href="https://www.techtarget.com/searchdatacenter/definition/kernel">https://www.techtarget.com/searchdatacenter/definition/kernel</a></p>

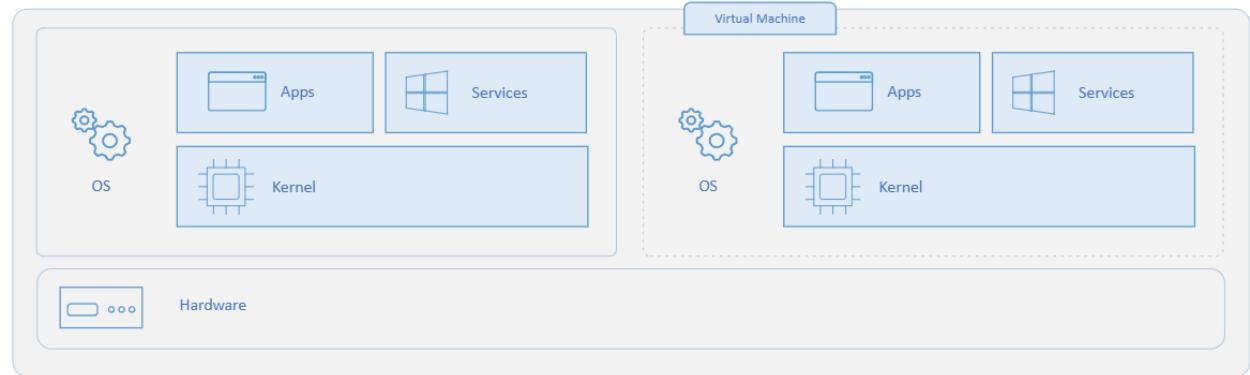
Claim 1	Accused Instrumentalities
<p>[1a] storing in memory accessible to at least some of the servers a plurality of secure containers of application software, each container comprising one or more of the executable applications and a set of associated system files required to execute the one or more applications, for use with a local kernel residing permanently on one of the servers;</p>	<p>The method practiced by Microsoft and/or its customer through the Accused Instrumentalities includes a step of storing in memory accessible to at least some of the servers a plurality of secure containers of application software, each container comprising one or more of the executable applications and a set of associated system files required to execute the one or more applications, for use with a local kernel residing permanently on one of the servers.</p> <p>For example, OCI and/or OKE stores application containers, sometimes called Docker containers, container images, Kubernetes containers, or Kubernetes pods, in persistent storage available to each node running the application. The terms “node” and “host” are both used to refer to the claimed server. The container might be in a format defined by the Open Container Initiative. This storage may be physically attached to the server or connected through any supported interconnect, including over a network. In addition to the application software, each container includes associated system files, including a Linux user space required to execute the application, for example including runtime linked libraries (e.g. .dll/.so files), configuration files, Windows services, etc. necessary for the application. For example, the container includes a base OS image provided by Microsoft or by a third party. The container is for use with the host kernel, for example because the application(s) and container libraries are linked against the Linux kernel, and the supported host operating systems also use the Linux kernel, which has a stable binary interface. In another example, the container is for use with the host kernel because the application(s) and Windows services within the container are compatible with the host Windows operating system and kernel.</p> <p>The containers are secure containers as claimed. For example, the data within an individual container is insulated from the effects of other containers except to the extent the container is specifically configured to allow other containers to modify its data, for example using a shared volume.</p> <p><i>See, e.g.:</i></p>

Claim 1	Accused Instrumentalities
	<p><b>Docker benefits</b></p> <p>When we use Docker, we immediately get access to the benefits containerization offer.</p> <p><b>Efficient hardware use</b></p> <p>Containers run without using a virtual machine (VM). As we learned, the container relies on the host kernel for functions such as file system, network management, process scheduling, and memory management.</p>  <p>Compared to a VM, we can see that a VM requires an OS installed to provide kernel functions to the running applications inside the VM. Keep in mind that the VM OS also requires disk space, memory, and CPU time. By removing the VM and the additional OS requirement, we can free resources on the host and use it for running other containers.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p>

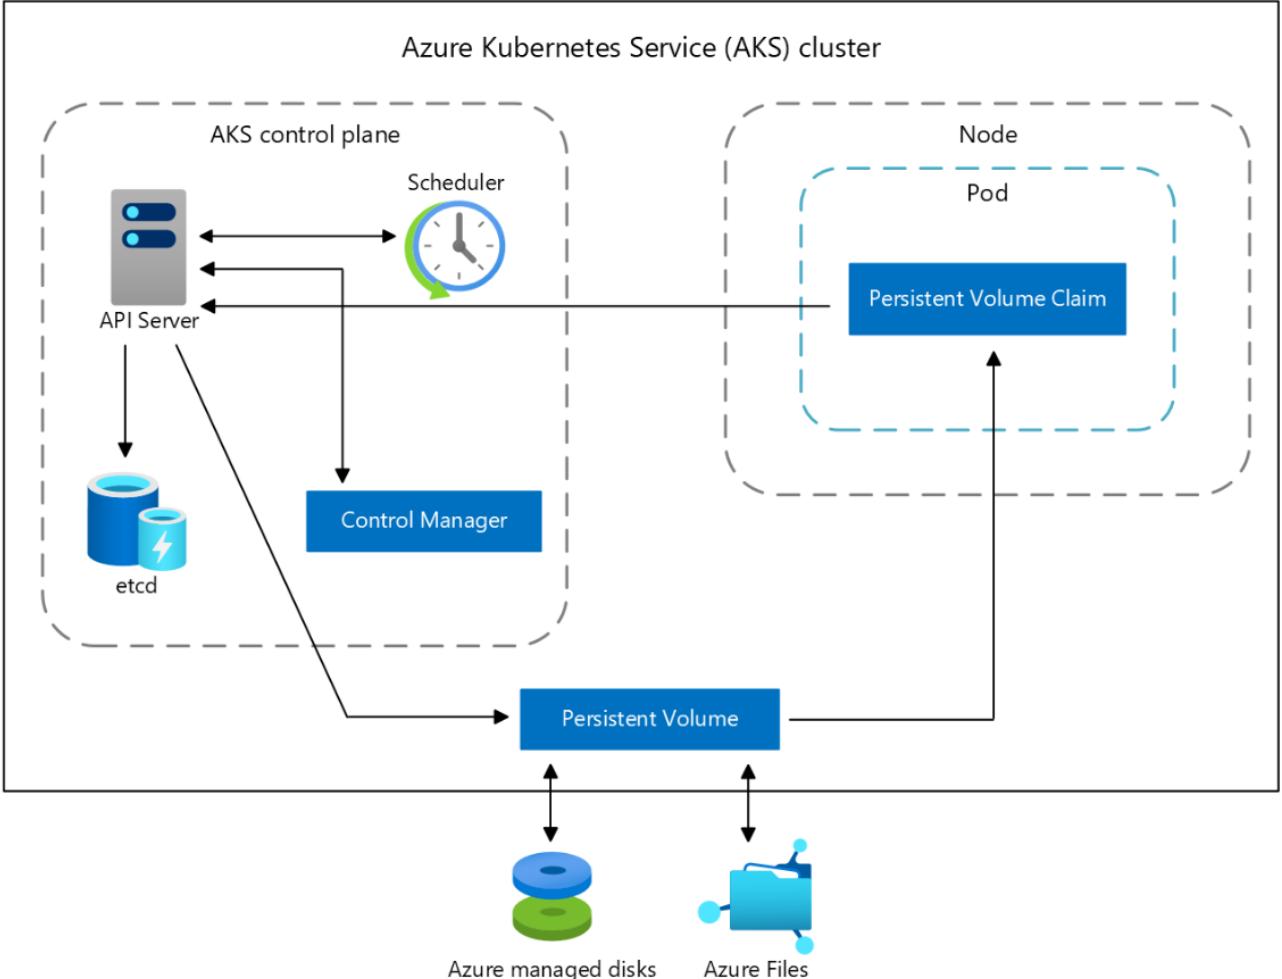
Claim 1	Accused Instrumentalities
	<p><b>Container isolation</b></p> <p>Docker containers provide security features to run multiple containers simultaneously on the same host without affecting each other. As we learned, we can configure both data storage and network configuration to isolate our containers or share data and connectivity between specific containers.</p> <p>Let's compare this feature to using VMs.</p>  <p>Assume we have a physical host running two VMs. We have three applications that we want to run isolated from each other. We decide to deploy the first app onto VM1 and the second onto VM2 to separate the two apps from each other. If we now choose to install the third application, we'll need to install another VM to continue this pattern.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p> <p><b>Application portability</b></p> <p>Containers run almost everywhere: desktops, physical servers, VMs, and in the cloud. This runtime compatibility makes it easy to move containerized applications among different environments.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p>

Claim 1	Accused Instrumentalities
	<p><b>Cloud deployments</b></p> <p>Docker containers are the default container architecture the Azure containerization services use, and many other cloud platforms also support them.</p> <p>For instance, you can deploy Docker containers to Azure Container Instances, Azure App Service, and Azure Kubernetes Services. Each of these options provides you with different features and capabilities.</p> <p>For example, Azure container instances allow you to focus on designing and building your applications without the overhead of managing infrastructure. When you have many containers to orchestrate, Azure Kubernetes service makes it easy to deploy and manage large-scale container deployments.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p> <ul style="list-style-type: none"> <li>• Deploy containers at scale on Azure or other clouds: <ul style="list-style-type: none"> <li>○ Pull your app (container image) from a container registry, such as the Azure Container Registry, and then deploy and manage it at scale using an orchestrator such as <a href="#">Azure Kubernetes Service (AKS)</a>.</li> <li>○ Azure Kubernetes Service deploys containers to Azure virtual machines and manages them at scale, whether that's dozens of containers, hundreds, or even thousands. The Azure virtual machines run either a customized Windows Server image (if you're deploying a Windows-based app), or a customized Ubuntu Linux image (if you're deploying a Linux-based app).</li> </ul> </li> </ul> <p><a href="https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/">https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/</a></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="688 225 1178 274">How containers work</h2> <p data-bbox="688 306 1867 409">A container is an isolated, lightweight package for running an application on the host operating system. Containers build on top of the host operating system's kernel (which can be thought of as the buried plumbing of the operating system), as shown in the diagram below.</p>  <p data-bbox="688 784 1867 964">While a container shares the host operating system's kernel, the container doesn't get unfettered access to it. Instead, the container gets an isolated—and in some cases virtualized—view of the system. For example, a container can access a virtualized version of the file system and registry, but any changes affect only the container and are discarded when it stops. To save data, the container can mount persistent storage such as an <a href="#">Azure Disk</a> or a file share (including <a href="#">Azure Files</a>).</p> <p data-bbox="688 997 1867 1209">A container builds on top of the kernel, but the kernel doesn't provide all of the APIs and services an app needs to run—most of these are provided by system files (libraries) that run above the kernel in user mode. Because a container is isolated from the host's user mode environment, the container needs its own copy of these user mode system files, which are packaged into something known as a base image. The base image serves as the foundational layer upon which your container is built, providing it with operating system services not provided by the kernel. But we'll talk more about container images later.</p> <p data-bbox="667 1233 1628 1266"><a href="https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/">https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/</a></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="688 221 1463 274">Containers vs. virtual machines</h2> <p data-bbox="688 311 1981 385">In contrast to a container, a virtual machine (VMs) runs a complete operating system—including its own kernel—as shown in this diagram.</p>  <p data-bbox="688 882 1938 988">Containers and virtual machines each have their uses—in fact, many deployments of containers use virtual machines as the host operating system rather than running directly on the hardware, especially when running containers in the cloud.</p> <p data-bbox="667 1013 1628 1046"><a href="https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/">https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/</a></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="677 202 1127 262">Container images</h2> <p data-bbox="677 295 1972 453">All containers are created from container images. A container image is a bundle of files organized into a stack of layers that resides on your local machine or in a remote container registry. The container image consists of the user mode operating system files needed to support your app, any runtimes or dependencies of your app, and any other miscellaneous configuration file your app needs to run properly.</p> <p data-bbox="677 486 1902 556">Microsoft offers several images (called base images) that you can use as a starting point to build your own container image:</p> <ul data-bbox="713 589 1972 829" style="list-style-type: none"> <li>• <b>Windows</b> - contains the full set of Windows APIs and system services (minus server roles).</li> <li>• <b>Windows Server</b> - contains the full set of Windows APIs and system services.</li> <li>• <b>Windows Server Core</b> - a smaller image that contains a subset of the Windows Server APIs—namely the full .NET framework. It also includes most but not all server roles (for example Fax Server is not included).</li> <li>• <b>Nano Server</b> - the smallest Windows Server image and includes support for the .NET Core APIs and some server roles.</li> </ul> <p data-bbox="677 866 1628 899"><a href="https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/">https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/</a></p> <p data-bbox="677 940 1972 1011">There are many different orchestrators that you can use with Windows containers; here are the options Microsoft provides:</p> <ul data-bbox="730 1046 1833 1116" style="list-style-type: none"> <li>• <a href="#">Azure Kubernetes Service (AKS)</a> - use a managed Azure Kubernetes service</li> <li>• <a href="#">Azure Kubernetes Service (AKS) on Azure Stack HCI</a> - use Azure Kubernetes Service on-premises</li> </ul> <p data-bbox="677 1170 1628 1202"><a href="https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/">https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/</a></p>

Claim 1	Accused Instrumentalities
	 <p>The diagram illustrates the architecture of an Azure Kubernetes Service (AKS) cluster. It is divided into two main sections: the AKS control plane and a Node. The AKS control plane (left) contains the API Server, etcd (represented by two cylinders), and a Control Manager. The Control Manager is connected to a Scheduler (represented by a clock icon). The Scheduler interacts with the API Server. The API Server is also connected to the Control Manager. The Control Manager is connected to a Persistent Volume. The Persistent Volume is connected to two storage options: Azure managed disks (represented by a blue and green cylinder) and Azure Files (represented by a blue folder icon). The Persistent Volume is also connected to a Pod, which is part of a Node (right). The Node is represented by a dashed box containing a Pod and a Persistent Volume Claim (represented by a blue box). The Persistent Volume Claim is connected to the Persistent Volume.</p> <p><a href="https://learn.microsoft.com/en-us/azure/aks/concepts-storage">https://learn.microsoft.com/en-us/azure/aks/concepts-storage</a></p>

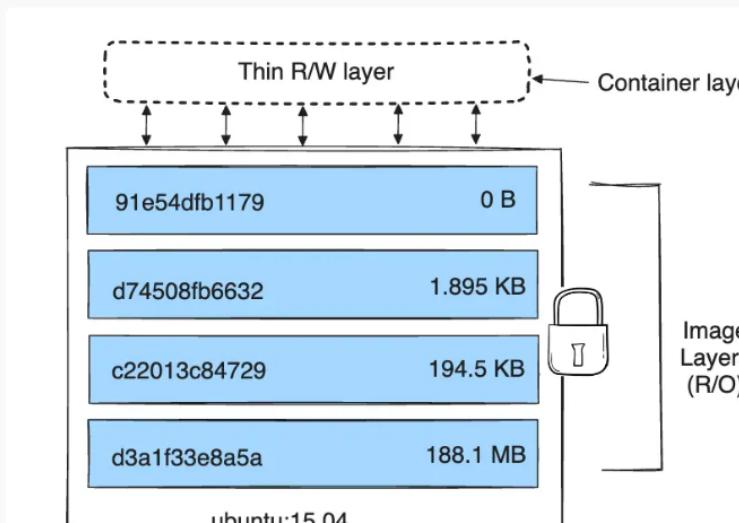
Claim 1	Accused Instrumentalities
	<p><b>Ephemeral OS disk</b></p> <p>By default, Azure automatically replicates the operating system disk for a virtual machine to Azure Storage to avoid data loss when the VM is relocated to another host. However, since containers aren't designed to have local state persisted, this behavior offers limited value while providing some drawbacks. These drawbacks include, but aren't limited to, slower node provisioning and higher read/write latency.</p> <p>By contrast, ephemeral OS disks are stored only on the host machine, just like a temporary disk. With this configuration, you get lower read/write latency, together with faster node scaling and cluster upgrades.</p> <p><a href="https://learn.microsoft.com/en-us/azure/aks/concepts-storage">https://learn.microsoft.com/en-us/azure/aks/concepts-storage</a></p> <p><b>Volumes</b></p> <p>Kubernetes typically treats individual pods as ephemeral, disposable resources. Applications have different approaches available to them for using and persisting data. A <i>volume</i> represents a way to store, retrieve, and persist data across pods and through the application lifecycle.</p> <p>Traditional volumes are created as Kubernetes resources backed by Azure Storage. You can manually create data volumes to be assigned to pods directly or have Kubernetes automatically create them. Data volumes can use: <a href="#">Azure Disk</a>, <a href="#">Azure Files</a>, <a href="#">Azure NetApp Files</a>, or <a href="#">Azure Blobs</a>.</p> <p><a href="https://learn.microsoft.com/en-us/azure/aks/concepts-storage">https://learn.microsoft.com/en-us/azure/aks/concepts-storage</a></p>

Claim 1	Accused Instrumentalities
	<h2>Persistent volumes</h2> <p>Volumes defined and created as part of the pod lifecycle only exist until you delete the pod. Pods often expect their storage to remain if a pod is rescheduled on a different host during a maintenance event, especially in StatefulSets. A <i>persistent volume</i> (PV) is a storage resource created and managed by the Kubernetes API that can exist beyond the lifetime of an individual pod.</p> <p>You can use the following Azure Storage services to provide the persistent volume:</p> <ul style="list-style-type: none"> <li>• Azure Disk</li> <li>• Azure Files</li> <li>• Azure Container Storage</li> </ul> <p>As noted in the <a href="#">Volumes</a> section, the choice of Azure Disks or Azure Files is often determined by the need for concurrent access to the data or the performance tier.</p> <p><a href="https://learn.microsoft.com/en-us/azure/aks/concepts-storage">https://learn.microsoft.com/en-us/azure/aks/concepts-storage</a></p>

Claim 1	Accused Instrumentalities
	<p>Container security protects the entire end-to-end pipeline from build to the application workloads running in Azure Kubernetes Service (AKS).</p> <p>The Secure Supply Chain includes the build environment and registry.</p> <p>Kubernetes includes security components, such as <i>pod security standards</i> and <i>Secrets</i>. Azure includes components like Active Directory, Microsoft Defender for Containers, Azure Policy, Azure Key Vault, network security groups, and orchestrated cluster upgrades. AKS combines these security components to:</p> <ul style="list-style-type: none"> <li>• Provide a complete authentication and authorization story.</li> <li>• Apply AKS Built-in Azure Policy to secure your applications.</li> <li>• End-to-End insight from build through your application with Microsoft Defender for Containers.</li> <li>• Keep your AKS cluster running the latest OS security updates and Kubernetes releases.</li> <li>• Provide secure pod traffic and access to sensitive credentials.</li> </ul> <p><a href="https://learn.microsoft.com/en-us/azure/aks/concepts-security">https://learn.microsoft.com/en-us/azure/aks/concepts-security</a></p> <h2>Container images</h2> <p>A <a href="#">container image</a> is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <p><a href="https://kubernetes.io/docs/concepts/containers/">https://kubernetes.io/docs/concepts/containers/</a></p> <p><b>6. Do Docker containers package up the entire OS and make it easier to deploy?</b></p> <p>Docker containers do not package up the OS. They package up the applications with everything that the application needs to run. The engine is installed on top of the OS running on a host. Containers share the OS kernel allowing a single host to run multiple containers.</p> <p><a href="https://www.docker.com/blog/the-10-most-common-questions-it-admins-ask-about-docker/">https://www.docker.com/blog/the-10-most-common-questions-it-admins-ask-about-docker/</a></p>

Claim 1	Accused Instrumentalities
	<h2>About storage drivers</h2> <p>To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h3>Storage drivers versus Docker volumes</h3> <p>Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p>Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the <a href="#">volumes section</a> to learn how to use volumes to persist data and improve performance.</p> <p><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>

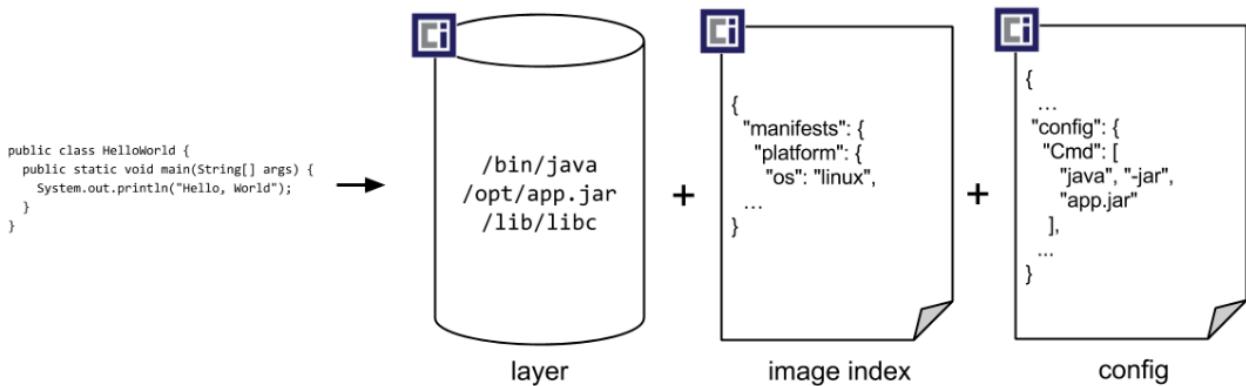
Claim 1	Accused Instrumentalities
	<h2 data-bbox="692 208 1136 262">Images and layers</h2> <p data-bbox="692 300 1877 376">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="713 437 1501 763"># syntax=docker/dockerfile:1  FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py</pre> <p data-bbox="692 829 1959 1139">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="692 1160 1284 1192"><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>

Claim 1	Accused Instrumentalities
	<p>Each layer is only a set of differences from the layer before it. Note that both <i>adding</i>, and <i>removing</i> files will result in a new layer. In the example above, the <code>\$HOME/.cache</code> directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the <a href="#">Best practices for writing Dockerfiles</a> and <a href="#">use multi-stage builds</a> sections to learn how to optimize your Dockerfiles for efficient images.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an <code>ubuntu:15.04</code> image.</p>  <p>Container (based on <code>ubuntu:15.04</code> image)</p> <p><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="688 221 967 283">Volumes</h2> <p data-bbox="688 339 1959 474">Volumes are the preferred mechanism for persisting data generated by and used by Docker containers. While <a href="#">bind mounts</a> are dependent on the directory structure and OS of the host machine, volumes are completely managed by Docker. Volumes have several advantages over bind mounts:</p> <p data-bbox="688 494 1353 523"><a href="https://kubernetes.io/docs/concepts/storage/volumes/">https://kubernetes.io/docs/concepts/storage/volumes/</a></p> <h2 data-bbox="688 576 1275 621">Container environment</h2> <p data-bbox="688 660 1522 729">The Kubernetes Container environment provides several important resources to Containers:</p> <ul data-bbox="734 768 1495 931" style="list-style-type: none"> <li data-bbox="734 768 1495 837">• A filesystem, which is a combination of an <a href="#">image</a> and one or more <a href="#">volumes</a>.</li> <li data-bbox="734 850 1248 878">• Information about the Container itself.</li> <li data-bbox="734 891 1332 918">• Information about other objects in the cluster.</li> </ul> <p data-bbox="688 967 1579 997"><a href="https://kubernetes.io/docs/concepts/containers/container-environment/">https://kubernetes.io/docs/concepts/containers/container-environment/</a></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="699 213 925 279">Images</h2> <p data-bbox="699 306 1575 463">A container image represents binary data that encapsulates an application and all its software dependencies. Container images are executable software bundles that can run standalone and that make very well defined assumptions about their runtime environment.</p> <p data-bbox="699 496 1584 572">You typically create a container image of your application and push it to a registry before referring to it in a Pod.</p> <p data-bbox="677 597 1374 633"><a href="https://kubernetes.io/docs/concepts/containers/images/">https://kubernetes.io/docs/concepts/containers/images/</a></p> <h2 data-bbox="699 675 967 740">Volumes</h2> <p data-bbox="699 776 1584 1225">On-disk files in a container are ephemeral, which presents some problems for non-trivial applications when running in containers. One problem occurs when a container crashes or is stopped. Container state is not saved so all of the files that were created or modified during the lifetime of the container are lost. During a crash, kubelet restarts the container with a clean state. Another problem occurs when multiple containers are running in a Pod and need to share files. It can be challenging to setup and access a shared filesystem across all of the containers. The Kubernetes volume abstraction solves both of these problems. Familiarity with Pods is suggested.</p> <p data-bbox="677 1250 1364 1286"><a href="https://kubernetes.io/docs/concepts/storage/volumes/">https://kubernetes.io/docs/concepts/storage/volumes/</a></p>

Claim 1	Accused Instrumentalities
	<p><b>Open Container Initiative</b></p> <hr/> <p><b>Image Format Specification</b></p> <hr/> <p>This specification defines an OCI Image, consisting of an <a href="#">image manifest</a>, an <a href="#">image index</a> (optional), a set of <a href="#">filesystem layers</a>, and a <a href="#">configuration</a>.</p> <p>The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p> <p><a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</a></p>

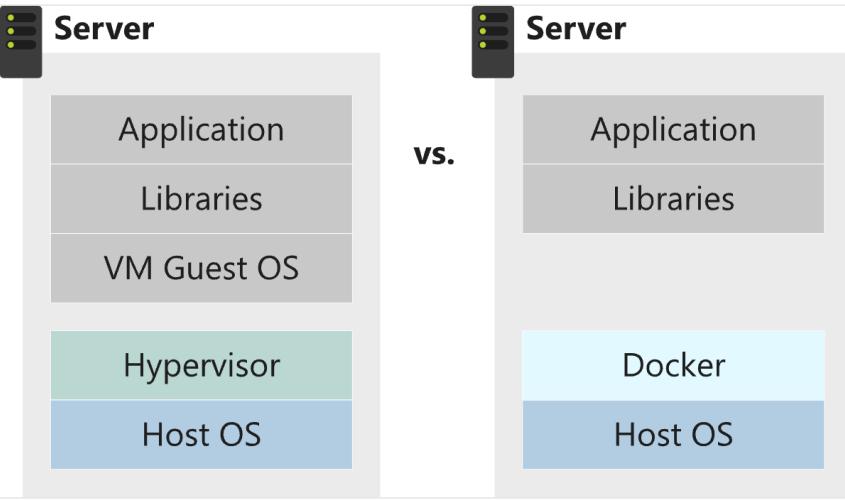
Claim 1	Accused Instrumentalities
	<h2>Overview</h2> <p>At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more <a href="#">filesystem layer changeset</a> archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p>  <pre data-bbox="713 731 1036 829"> public class HelloWorld {     public static void main(String[] args) {         System.out.println("Hello, World");     } } </pre> <p><a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</a></p>

Claim 1	Accused Instrumentalities
	<h2>OCI Image Configuration</h2> <p>An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in <a href="#">Layers</a>.</p> <p>This section defines the <code>application/vnd.oci.image.config.v1+json</code> <a href="#">media type</a>.</p> <p><a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</a></p>

Claim 1	Accused Instrumentalities
	<p><b>Layer</b></p> <ul style="list-style-type: none"> <li>• Image filesystems are composed of <i>layers</i>.</li> <li>• Each layer represents a set of filesystem changes in a tar-based <a href="#">layer format</a>, recording files to be added, changed, or deleted relative to its parent layer.</li> <li>• Layers do not have configuration metadata such as environment variables or default arguments - these are properties of the image as a whole rather than any particular layer.</li> <li>• Using a layer-based or union filesystem such as AUFS, or by computing the diff from filesystem snapshots, the filesystem changeset can be used to present a series of image layers as if they were one cohesive filesystem.</li> </ul> <p><b>Image JSON</b></p> <ul style="list-style-type: none"> <li>• Each image has an associated JSON structure which describes some basic information about the image such as date created, author, as well as execution/runtime configuration like its entrypoint, default arguments, networking, and volumes.</li> <li>• The JSON structure also references a cryptographic hash of each layer used by the image, and provides history information for those layers.</li> <li>• This JSON is considered to be immutable, because changing it would change the computed <a href="#">ImageID</a>.</li> <li>• Changing it means creating a new derived image, instead of changing the existing image.</li> </ul> <p><a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</a></p>

Claim 1	Accused Instrumentalities
	<ul style="list-style-type: none"> <li>• <b>rootfs</b> <i>object</i>, REQUIRED</li> </ul> <p>The rootfs key references the layer content addresses used by the image. This makes the image config hash depend on the filesystem hash.</p> <ul style="list-style-type: none"> <li>◦ <b>type</b> <i>string</i>, REQUIRED</li> </ul> <p>MUST be set to <code>layers</code>. Implementations MUST generate an error if they encounter an unknown value while verifying or unpacking an image.</p> <ul style="list-style-type: none"> <li>◦ <b>diff_ids</b> <i>array of strings</i>, REQUIRED</li> </ul> <p>An array of layer content hashes (<code>DiffIDs</code>), in order from first to last.</p> <p><a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</a></p>

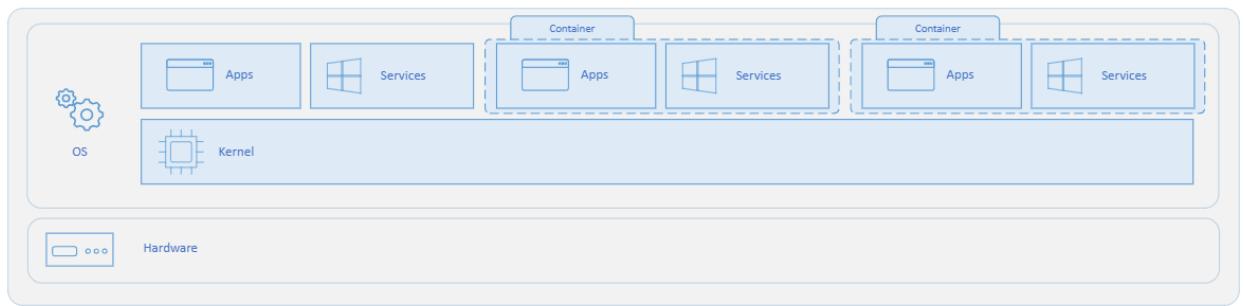
Claim 1	Accused Instrumentalities
<p>[1b] wherein the set of associated system files are compatible with a local kernel of at least some of the plurality of different operating systems,</p>	<p>In the method practiced by Microsoft and/or its customer through the Accused Instrumentalities, the set of associated system files are compatible with a local kernel of at least some of the plurality of different operating systems.</p> <p>The associated system files in the container are compatible with the host kernel, for example because they are linked against the Linux kernel and the supported host operating systems also use the Linux kernel, which has a stable binary interface. In another example, the associated system files are linked against and compatible with a Windows operating system and kernel, and the host operating system runs a compatible Windows operating system with compatible Windows kernel.</p> <p><i>See discussion in element [1a] above.</i></p> <p><i>See also, e.g.:</i></p> <p><b>OS</b></p> <p>AKS supports Ubuntu 22.04 and Azure Linux 2.0 as the node OS for Linux node pools. For Windows node pools, AKS supports Windows Server 2022 as the default OS. Windows Server 2019 is being retired after Kubernetes version 1.32 reaches end of life and isn't supported in future releases. If you need to upgrade your Windows OS version, see <a href="#">Upgrade from Windows Server 2019 to Windows Server 2022</a>. For more information on using Windows Server on AKS, see <a href="#">Windows container considerations in Azure Kubernetes Service (AKS)</a>.</p> <p><a href="https://learn.microsoft.com/en-us/azure/aks/core-aks-concepts">https://learn.microsoft.com/en-us/azure/aks/core-aks-concepts</a></p> <p><b>Container runtime</b></p> <p>A container runtime is software that executes containers and manages container images on a node. The runtime helps abstract away sys-calls or OS-specific functionality to run containers on Linux or Windows. For Linux node pools, <a href="#">containerd</a> is used on Kubernetes version 1.19 and higher. For Windows Server 2019 and 2022 node pools, <a href="#">containerd</a> is generally available and is the only runtime option on Kubernetes version 1.23 and higher.</p> <p><a href="https://learn.microsoft.com/en-us/azure/aks/core-aks-concepts">https://learn.microsoft.com/en-us/azure/aks/core-aks-concepts</a></p>

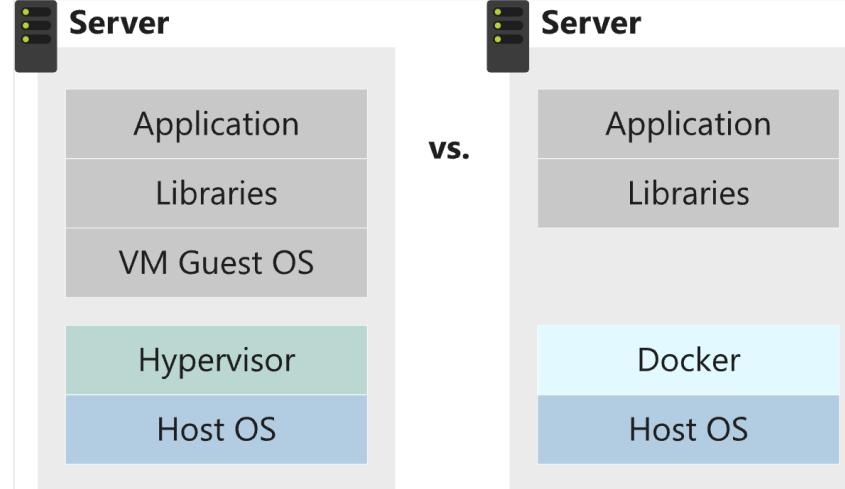
Claim 1	Accused Instrumentalities
	<p><b>Docker benefits</b></p> <p>When we use Docker, we immediately get access to the benefits containerization offer.</p> <p><b>Efficient hardware use</b></p> <p>Containers run without using a virtual machine (VM). As we learned, the container relies on the host kernel for functions such as file system, network management, process scheduling, and memory management.</p>  <p>Compared to a VM, we can see that a VM requires an OS installed to provide kernel functions to the running applications inside the VM. Keep in mind that the VM OS also requires disk space, memory, and CPU time. By removing the VM and the additional OS requirement, we can free resources on the host and use it for running other containers.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p>

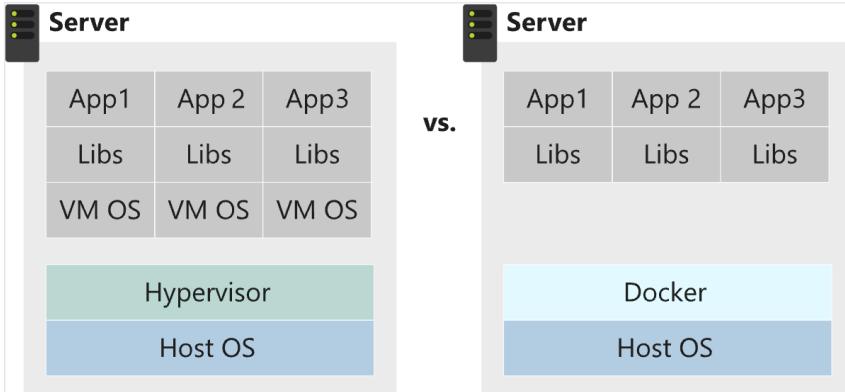
Claim 1	Accused Instrumentalities
	<p><b>Application portability</b></p> <p>Containers run almost everywhere: desktops, physical servers, VMs, and in the cloud. This runtime compatibility makes it easy to move containerized applications among different environments.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p> <p><b>Cloud deployments</b></p> <p>Docker containers are the default container architecture the Azure containerization services use, and many other cloud platforms also support them.</p> <p>For instance, you can deploy Docker containers to Azure Container Instances, Azure App Service, and Azure Kubernetes Services. Each of these options provides you with different features and capabilities.</p> <p>For example, Azure container instances allow you to focus on designing and building your applications without the overhead of managing infrastructure. When you have many containers to orchestrate, Azure Kubernetes service makes it easy to deploy and manage large-scale container deployments.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p>

Claim 1	Accused Instrumentalities
	<h2>Windows nodes in Kubernetes</h2> <p>To enable the orchestration of Windows containers in Kubernetes, include Windows nodes in your existing Linux cluster. Scheduling Windows containers in Pods on Kubernetes is similar to scheduling Linux-based containers.</p> <p>In order to run Windows containers, your Kubernetes cluster must include multiple operating systems. While you can only run the control plane on Linux, you can deploy worker nodes running either Windows or Linux.</p> <p>Windows nodes are supported provided that the operating system is Windows Server 2019 or Windows Server 2022.</p> <p>This document uses the term <i>Windows containers</i> to mean Windows containers with process isolation. Kubernetes does not support running Windows containers with Hyper-V isolation.</p> <p><a href="https://kubernetes.io/docs/concepts/windows/intro/">https://kubernetes.io/docs/concepts/windows/intro/</a></p>

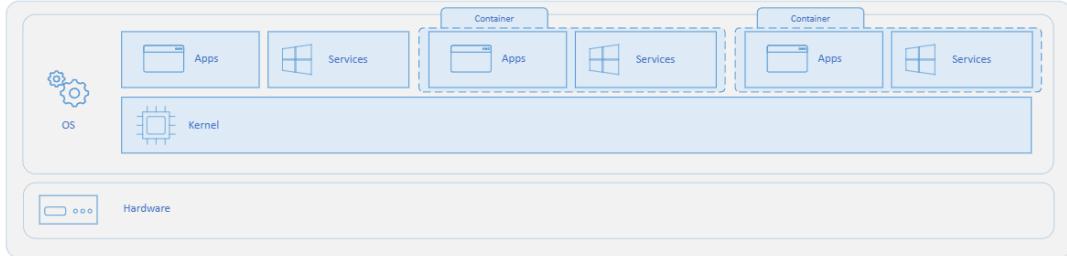
Claim 1	Accused Instrumentalities
	<h2>Windows OS version compatibility</h2> <p>On Windows nodes, strict compatibility rules apply where the host OS version must match the container base image OS version. Only Windows containers with a container operating system of Windows Server 2019 are fully supported.</p> <p>For Kubernetes v1.31, operating system compatibility for Windows nodes (and Pods) is as follows:</p> <p><b>Windows Server LTSC release</b></p> <ul style="list-style-type: none"><li>Windows Server 2019</li><li>Windows Server 2022</li></ul> <p><b>Windows Server SAC release</b></p> <ul style="list-style-type: none"><li>Windows Server version 20H2</li></ul> <p>The Kubernetes <a href="#">version-skew policy</a> also applies.</p> <p><a href="https://kubernetes.io/docs/concepts/windows/intro/#windows-os-version-support">https://kubernetes.io/docs/concepts/windows/intro/#windows-os-version-support</a></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="688 230 1227 279">How containers work</h2> <p data-bbox="688 319 1981 434">A container is an isolated, lightweight package for running an application on the host operating system. Containers build on top of the host operating system's kernel (which can be thought of as the buried plumbing of the operating system), as shown in the diagram below.</p>  <p data-bbox="688 842 1981 1037">While a container shares the host operating system's kernel, the container doesn't get unfettered access to it. Instead, the container gets an isolated—and in some cases virtualized—view of the system. For example, a container can access a virtualized version of the file system and registry, but any changes affect only the container and are discarded when it stops. To save data, the container can mount persistent storage such as an <a href="#">Azure Disk</a> or a file share (including <a href="#">Azure Files</a>).</p> <p data-bbox="688 1070 1981 1307">A container builds on top of the kernel, but the kernel doesn't provide all of the APIs and services an app needs to run—most of these are provided by system files (libraries) that run above the kernel in user mode. Because a container is isolated from the host's user mode environment, the container needs its own copy of these user mode system files, which are packaged into something known as a base image. The base image serves as the foundational layer upon which your container is built, providing it with operating system services not provided by the kernel. But we'll talk more about container images later.</p> <p data-bbox="671 1326 1628 1359"><a href="https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/">https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/</a></p>

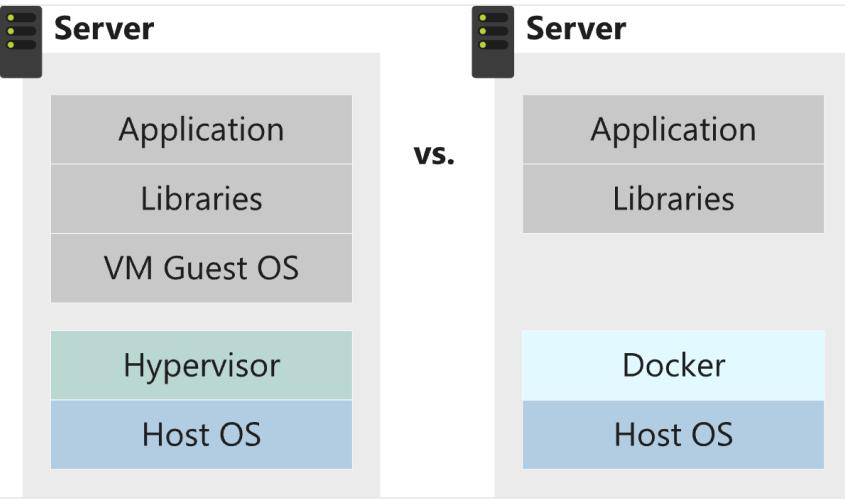
Claim 1	Accused Instrumentalities
<p>[1c] the containers of application software excluding a kernel,</p> <p><i>See, e.g.:</i></p> <p><b>Docker benefits</b></p> <p>When we use Docker, we immediately get access to the benefits containerization offer.</p> <p><b>Efficient hardware use</b></p> <p>Containers run without using a virtual machine (VM). As we learned, the container relies on the host kernel for functions such as file system, network management, process scheduling, and memory management.</p>  <p>Compared to a VM, we can see that a VM requires an OS installed to provide kernel functions to the running applications inside the VM. Keep in mind that the VM OS also requires disk space, memory, and CPU time. By removing the VM and the additional OS requirement, we can free resources on the host and use it for running other containers.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p>	

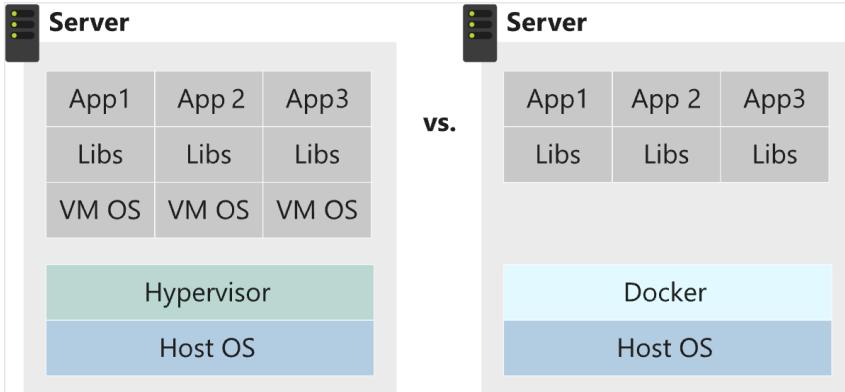
Claim 1	Accused Instrumentalities
	<p><b>Container isolation</b></p> <p>Docker containers provide security features to run multiple containers simultaneously on the same host without affecting each other. As we learned, we can configure both data storage and network configuration to isolate our containers or share data and connectivity between specific containers.</p> <p>Let's compare this feature to using VMs.</p>  <p>Assume we have a physical host running two VMs. We have three applications that we want to run isolated from each other. We decide to deploy the first app onto VM1 and the second onto VM2 to separate the two apps from each other. If we now choose to install the third application, we'll need to install another VM to continue this pattern.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p> <p><b>Application portability</b></p> <p>Containers run almost everywhere: desktops, physical servers, VMs, and in the cloud. This runtime compatibility makes it easy to move containerized applications among different environments.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p>

Claim 1	Accused Instrumentalities
	<p><b>Cloud deployments</b></p> <p>Docker containers are the default container architecture the Azure containerization services use, and many other cloud platforms also support them.</p> <p>For instance, you can deploy Docker containers to Azure Container Instances, Azure App Service, and Azure Kubernetes Services. Each of these options provides you with different features and capabilities.</p> <p>For example, Azure container instances allow you to focus on designing and building your applications without the overhead of managing infrastructure. When you have many containers to orchestrate, Azure Kubernetes service makes it easy to deploy and manage large-scale container deployments.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p>

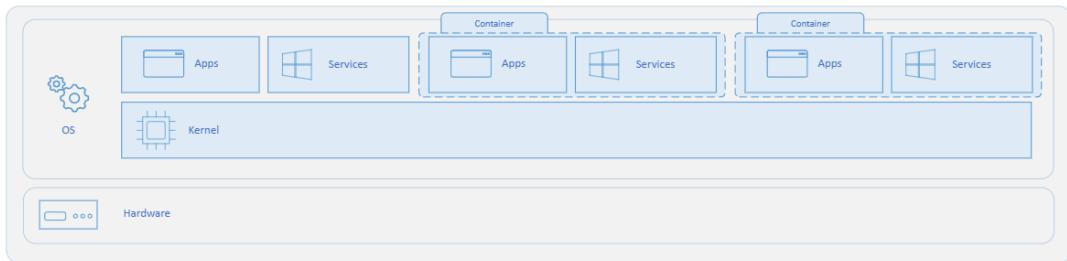
Claim 1	Accused Instrumentalities
	<h2 data-bbox="688 225 1157 269">How containers work</h2> <p data-bbox="688 303 1803 401">A container is an isolated, lightweight package for running an application on the host operating system. Containers build on top of the host operating system's kernel (which can be thought of as the buried plumbing of the operating system), as shown in the diagram below.</p>  <p data-bbox="688 752 1803 915">While a container shares the host operating system's kernel, the container doesn't get unfettered access to it. Instead, the container gets an isolated—and in some cases virtualized—view of the system. For example, a container can access a virtualized version of the file system and registry, but any changes affect only the container and are discarded when it stops. To save data, the container can mount persistent storage such as an <a href="#">Azure Disk</a> or a file share (including <a href="#">Azure Files</a>).</p> <p data-bbox="688 948 1803 1152">A container builds on top of the kernel, but the kernel doesn't provide all of the APIs and services an app needs to run—most of these are provided by system files (libraries) that run above the kernel in user mode. Because a container is isolated from the host's user mode environment, the container needs its own copy of these user mode system files, which are packaged into something known as a base image. The base image serves as the foundational layer upon which your container is built, providing it with operating system services not provided by the kernel. But we'll talk more about container images later.</p> <p data-bbox="688 1176 1628 1209"><a href="https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/">https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/</a></p> <p data-bbox="688 1241 1459 1266"><b>6. Do Docker containers package up the entire OS and make it easier to deploy?</b></p> <p data-bbox="688 1299 1723 1388">Docker containers do not package up the OS. They package up the applications with everything that the application needs to run. The engine is installed on top of the OS running on a host. Containers share the OS kernel allowing a single host to run multiple containers.</p> <p data-bbox="688 1396 1856 1429"><a href="https://www.docker.com/blog/the-10-most-common-questions-it-admins-ask-about-docker/">https://www.docker.com/blog/the-10-most-common-questions-it-admins-ask-about-docker/</a></p>

Claim 1	Accused Instrumentalities
<p>[1d] wherein some or all of the associated system files within a container stored in memory are utilized in place of the associated local system files that remain resident on the server,</p>	<p>In the method practiced by Microsoft and/or its customer through the Accused Instrumentalities, some or all of the associated system files within a container stored in memory are utilized in place of the associated local system files that remain resident on the server.</p> <p>For example, each container will utilize its own associated system files, including runtime linked libraries (e.g. .dll/.so files), configuration files, etc., not the corresponding associated local system files (e.g., libraries and configuration files of the host OS). As described above and below, in the Accused Instrumentalities the associated system files provide at least some of the same functionalities as the associated local system files. The host/node's associated local system files remain resident on the host/node, for example for use by system processes or applications outside the container environment.</p> <p><i>See, e.g.:</i></p>

Claim 1	Accused Instrumentalities
	<p><b>Docker benefits</b></p> <p>When we use Docker, we immediately get access to the benefits containerization offer.</p> <p><b>Efficient hardware use</b></p> <p>Containers run without using a virtual machine (VM). As we learned, the container relies on the host kernel for functions such as file system, network management, process scheduling, and memory management.</p>  <p>Compared to a VM, we can see that a VM requires an OS installed to provide kernel functions to the running applications inside the VM. Keep in mind that the VM OS also requires disk space, memory, and CPU time. By removing the VM and the additional OS requirement, we can free resources on the host and use it for running other containers.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p>

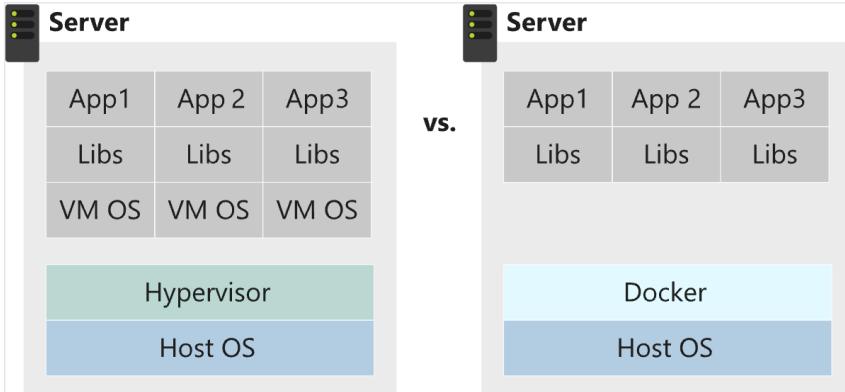
Claim 1	Accused Instrumentalities
	<p><b>Container isolation</b></p> <p>Docker containers provide security features to run multiple containers simultaneously on the same host without affecting each other. As we learned, we can configure both data storage and network configuration to isolate our containers or share data and connectivity between specific containers.</p> <p>Let's compare this feature to using VMs.</p>  <p>Assume we have a physical host running two VMs. We have three applications that we want to run isolated from each other. We decide to deploy the first app onto VM1 and the second onto VM2 to separate the two apps from each other. If we now choose to install the third application, we'll need to install another VM to continue this pattern.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p> <p><b>Application portability</b></p> <p>Containers run almost everywhere: desktops, physical servers, VMs, and in the cloud. This runtime compatibility makes it easy to move containerized applications among different environments.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p>

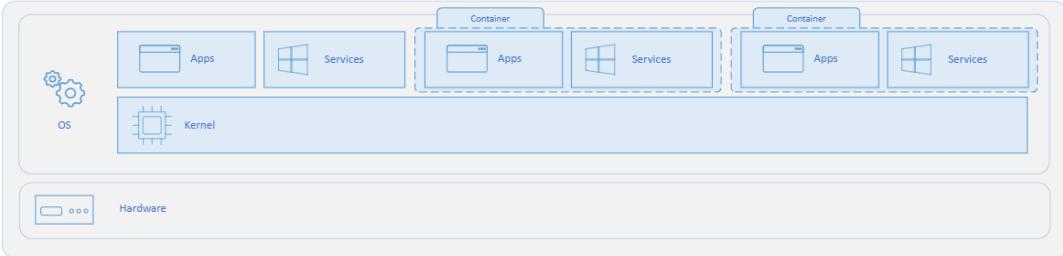
Claim 1	Accused Instrumentalities
	<p><b>Cloud deployments</b></p> <p>Docker containers are the default container architecture the Azure containerization services use, and many other cloud platforms also support them.</p> <p>For instance, you can deploy Docker containers to Azure Container Instances, Azure App Service, and Azure Kubernetes Services. Each of these options provides you with different features and capabilities.</p> <p>For example, Azure container instances allow you to focus on designing and building your applications without the overhead of managing infrastructure. When you have many containers to orchestrate, Azure Kubernetes service makes it easy to deploy and manage large-scale container deployments.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p> <p><b>Container runtime</b></p> <p>A container runtime is software that executes containers and manages container images on a node. The runtime helps abstract away sys-calls or OS-specific functionality to run containers on Linux or Windows. For Linux node pools, <a href="#">containerd</a> is used on Kubernetes version 1.19 and higher. For Windows Server 2019 and 2022 node pools, <a href="#">containerd</a> is generally available and is the only runtime option on Kubernetes version 1.23 and higher.</p> <p><a href="https://learn.microsoft.com/en-us/azure/aks/core-aks-concepts">https://learn.microsoft.com/en-us/azure/aks/core-aks-concepts</a></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="688 225 1157 269">How containers work</h2> <p data-bbox="688 303 1803 401">A container is an isolated, lightweight package for running an application on the host operating system. Containers build on top of the host operating system's kernel (which can be thought of as the buried plumbing of the operating system), as shown in the diagram below.</p>  <p data-bbox="688 752 1803 915">While a container shares the host operating system's kernel, the container doesn't get unfettered access to it. Instead, the container gets an isolated—and in some cases virtualized—view of the system. For example, a container can access a virtualized version of the file system and registry, but any changes affect only the container and are discarded when it stops. To save data, the container can mount persistent storage such as an <a href="#">Azure Disk</a> or a file share (including <a href="#">Azure Files</a>).</p> <p data-bbox="688 948 1803 1152">A container builds on top of the kernel, but the kernel doesn't provide all of the APIs and services an app needs to run—most of these are provided by system files (libraries) that run above the kernel in user mode. Because a container is isolated from the host's user mode environment, the container needs its own copy of these user mode system files, which are packaged into something known as a base image. The base image serves as the foundational layer upon which your container is built, providing it with operating system services not provided by the kernel. But we'll talk more about container images later.</p> <p data-bbox="688 1176 1628 1209"><a href="https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/">https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/</a></p> <p data-bbox="688 1237 1465 1263"><b>6. Do Docker containers package up the entire OS and make it easier to deploy?</b></p> <p data-bbox="688 1295 1723 1393">Docker containers do not package up the OS. They package up the applications with everything that the application needs to run. The engine is installed on top of the OS running on a host. Containers share the OS kernel allowing a single host to run multiple containers.</p> <p data-bbox="688 1400 1856 1432"><a href="https://www.docker.com/blog/the-10-most-common-questions-it-admins-ask-about-docker/">https://www.docker.com/blog/the-10-most-common-questions-it-admins-ask-about-docker/</a></p>

<b>Claim 1</b>	<b>Accused Instrumentalities</b>
<p>[1e] wherein said associated system files utilized in place of the associated local system files are copies or modified copies of the associated local system files that remain resident on the server,</p>	<p>In the method practiced by Microsoft and/or its customer through the Accused Instrumentalities, said associated system files utilized in place of the associated local system files are copies or modified copies of the associated local system files that remain resident on the server.</p> <p>For example, in some cases the host OS and container will use one or more identical system files, for example when both the host and the container incorporate the same Linux distribution version, or when both host and container use the same Windows services. In the case where the associated system files are identical to the associated local system files, they are copies thereof. In other cases modified copies are used instead, for example when different versions of the same library, or configuration files with different parameters, are used by the host and container.</p> <p><i>See discussion in elements [1a], [1b] above.</i></p>

Claim 1	Accused Instrumentalities
<p>[1f] and wherein the application software cannot be shared between the plurality of secure containers of application software,</p>	<p>In the method practiced by Microsoft and/or its customer through the Accused Instrumentalities, the application software cannot be shared between the plurality of secure containers of application software.</p> <p>For example, each container has an isolated runtime environment that cannot be accessed by other containers, for example including a per-container writeable layer or other ephemeral per-container storage. For another example, when the plurality of secure containers each corresponds to a different container image, each container cannot access another container's image and therefore application software.</p> <p><i>See, e.g.:</i></p>

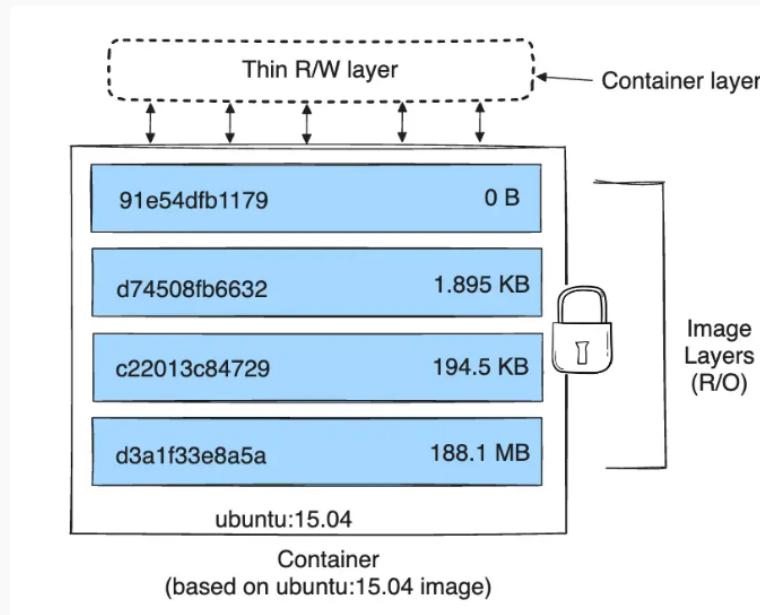
Claim 1	Accused Instrumentalities
	<p><b>Container isolation</b></p> <p>Docker containers provide security features to run multiple containers simultaneously on the same host without affecting each other. As we learned, we can configure both data storage and network configuration to isolate our containers or share data and connectivity between specific containers.</p> <p>Let's compare this feature to using VMs.</p>  <p>Assume we have a physical host running two VMs. We have three applications that we want to run isolated from each other. We decide to deploy the first app onto VM1 and the second onto VM2 to separate the two apps from each other. If we now choose to install the third application, we'll need to install another VM to continue this pattern.</p> <p><a href="https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers">https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/5-when-use-docker-containers</a></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="692 225 1157 269">How containers work</h2> <p data-bbox="692 303 1803 401">A container is an isolated, lightweight package for running an application on the host operating system. Containers build on top of the host operating system's kernel (which can be thought of as the buried plumbing of the operating system), as shown in the diagram below.</p>  <p data-bbox="692 752 1803 915">While a container shares the host operating system's kernel, the container doesn't get unfettered access to it. Instead, the container gets an isolated—and in some cases virtualized—view of the system. For example, a container can access a virtualized version of the file system and registry, but any changes affect only the container and are discarded when it stops. To save data, the container can mount persistent storage such as an <a href="#">Azure Disk</a> or a file share (including <a href="#">Azure Files</a>).</p> <p data-bbox="692 948 1803 1155">A container builds on top of the kernel, but the kernel doesn't provide all of the APIs and services an app needs to run—most of these are provided by system files (libraries) that run above the kernel in user mode. Because a container is isolated from the host's user mode environment, the container needs its own copy of these user mode system files, which are packaged into something known as a base image. The base image serves as the foundational layer upon which your container is built, providing it with operating system services not provided by the kernel. But we'll talk more about container images later.</p> <p data-bbox="677 1179 1628 1212"><a href="https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/">https://learn.microsoft.com/en-us/virtualization/windowscontainers/about/</a></p>

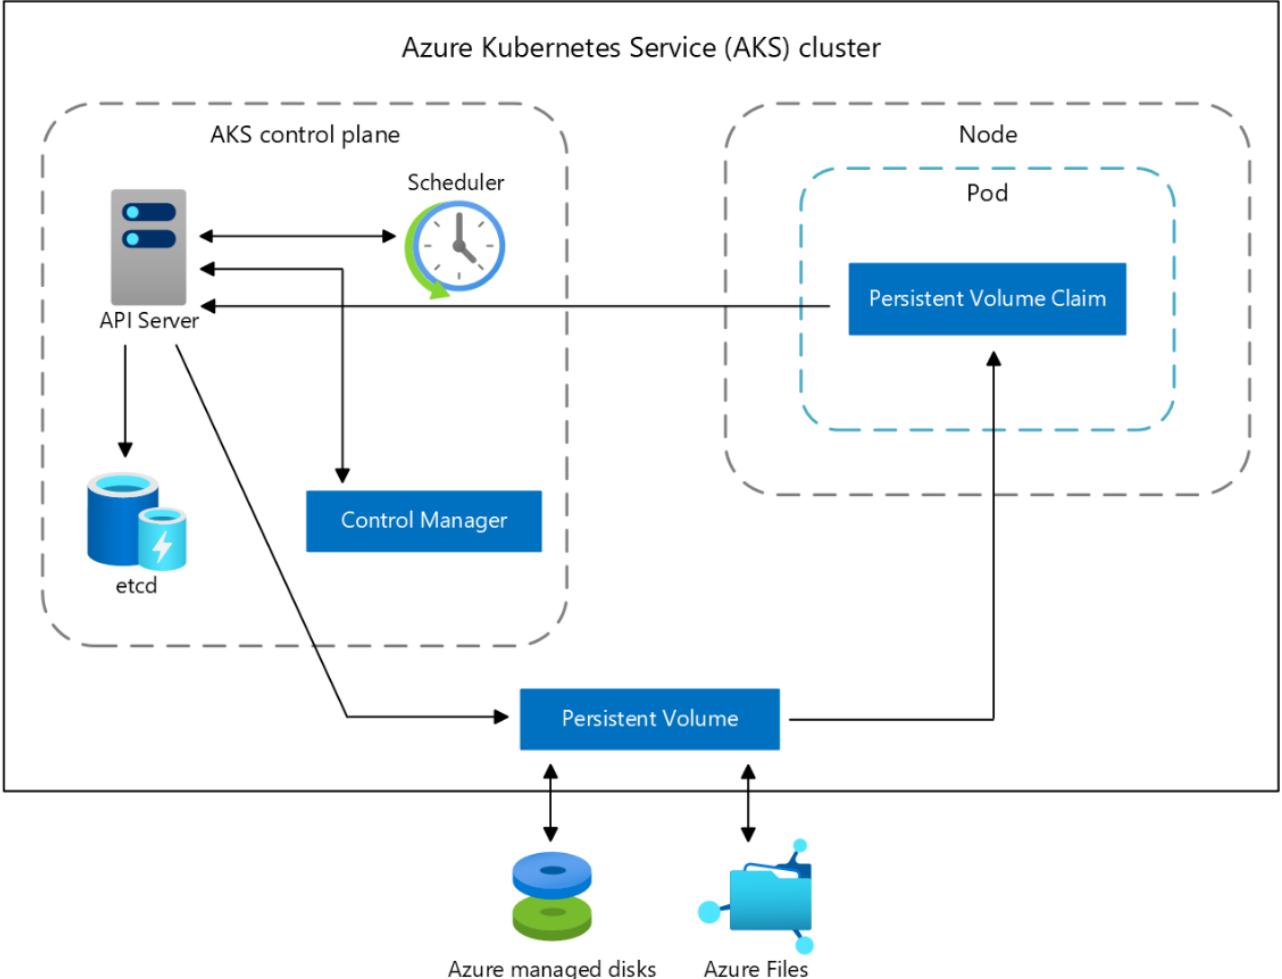
Claim 1	Accused Instrumentalities
	<p>Container security protects the entire end-to-end pipeline from build to the application workloads running in Azure Kubernetes Service (AKS).</p> <p>The Secure Supply Chain includes the build environment and registry.</p> <p>Kubernetes includes security components, such as <i>pod security standards</i> and <i>Secrets</i>. Azure includes components like Active Directory, Microsoft Defender for Containers, Azure Policy, Azure Key Vault, network security groups, and orchestrated cluster upgrades. AKS combines these security components to:</p> <ul style="list-style-type: none"> <li>• Provide a complete authentication and authorization story.</li> <li>• Apply AKS Built-in Azure Policy to secure your applications.</li> <li>• End-to-End insight from build through your application with Microsoft Defender for Containers.</li> <li>• Keep your AKS cluster running the latest OS security updates and Kubernetes releases.</li> <li>• Provide secure pod traffic and access to sensitive credentials.</li> </ul> <p><a href="https://learn.microsoft.com/en-us/azure/aks/concepts-security">https://learn.microsoft.com/en-us/azure/aks/concepts-security</a></p>

Claim 1	Accused Instrumentalities
	<h2>About storage drivers</h2> <p>To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2>Storage drivers versus Docker volumes</h2> <p>Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p>Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the <a href="#">volumes section</a> to learn how to use volumes to persist data and improve performance.</p> <p><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="692 208 1136 264">Images and layers</h2> <p data-bbox="692 300 1881 376">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="713 437 1501 763"># syntax=docker/dockerfile:1  FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py</pre> <p data-bbox="692 829 1959 1139">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="671 1160 1284 1192"><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>

Claim 1	Accused Instrumentalities
	<p>Each layer is only a set of differences from the layer before it. Note that both <i>adding</i>, and <i>removing</i> files will result in a new layer. In the example above, the <code>\$HOME/.cache</code> directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the <a href="#">Best practices for writing Dockerfiles</a> and <a href="#">use multi-stage builds</a> sections to learn how to optimize your Dockerfiles for efficient images.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an <code>ubuntu:15.04</code> image.</p>  <p><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>

Claim 1	Accused Instrumentalities
<p>[1g] and wherein each of the containers has a unique root file system that is different from an operating system's root file system.</p> <p><i>See, e.g.:</i></p>	<p>In the method practiced by Microsoft and/or its customer through the Accused Instrumentalities, each of the containers has a unique root file system that is different from an operating system's root file system.</p> <p>For example, the container's root file system comprises the image layer(s), an ephemeral writeable layer (e.g., in Docker terminology the container layer), and optionally one or more volumes. This root file system is distinct and isolated from the host operating system's root file system.</p>

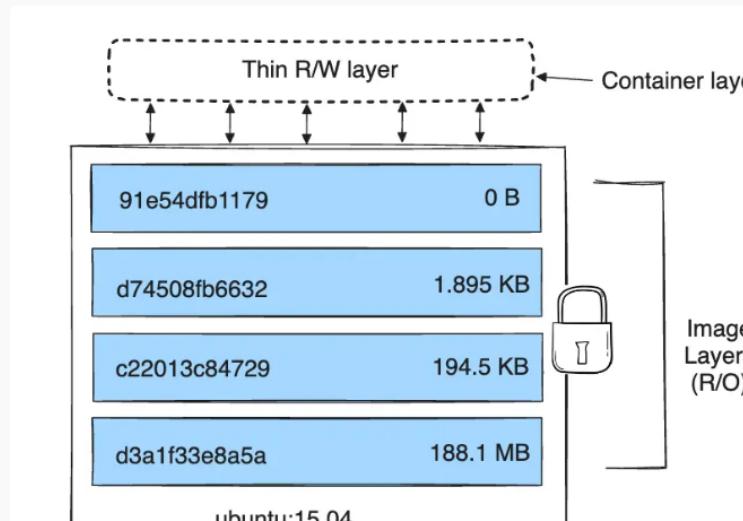
Claim 1	Accused Instrumentalities
	 <p>The diagram illustrates the architecture of an Azure Kubernetes Service (AKS) cluster. It is divided into two main sections: the AKS control plane and a Node. The AKS control plane (left) contains the API Server, etcd (represented by two cylinders), and a Control Manager. The Control Manager is connected to a Scheduler (represented by a clock icon). The Scheduler interacts with the API Server. The API Server is also connected to the etcd. The Control Manager is connected to a Persistent Volume. The Persistent Volume is connected to two storage options: Azure managed disks (represented by a blue and green cylinder) and Azure Files (represented by a blue folder icon with a blue dot). The Node (right) contains a Pod, which is associated with a Persistent Volume Claim. The Persistent Volume Claim is connected to the Persistent Volume.</p> <p><a href="https://learn.microsoft.com/en-us/azure/aks/concepts-storage">https://learn.microsoft.com/en-us/azure/aks/concepts-storage</a></p>

Claim 1	Accused Instrumentalities
	<p><b>Ephemeral OS disk</b></p> <p>By default, Azure automatically replicates the operating system disk for a virtual machine to Azure Storage to avoid data loss when the VM is relocated to another host. However, since containers aren't designed to have local state persisted, this behavior offers limited value while providing some drawbacks. These drawbacks include, but aren't limited to, slower node provisioning and higher read/write latency.</p> <p>By contrast, ephemeral OS disks are stored only on the host machine, just like a temporary disk. With this configuration, you get lower read/write latency, together with faster node scaling and cluster upgrades.</p> <p><a href="https://learn.microsoft.com/en-us/azure/aks/concepts-storage">https://learn.microsoft.com/en-us/azure/aks/concepts-storage</a></p> <p><b>Volumes</b></p> <p>Kubernetes typically treats individual pods as ephemeral, disposable resources. Applications have different approaches available to them for using and persisting data. A <i>volume</i> represents a way to store, retrieve, and persist data across pods and through the application lifecycle.</p> <p>Traditional volumes are created as Kubernetes resources backed by Azure Storage. You can manually create data volumes to be assigned to pods directly or have Kubernetes automatically create them. Data volumes can use: <a href="#">Azure Disk</a>, <a href="#">Azure Files</a>, <a href="#">Azure NetApp Files</a>, or <a href="#">Azure Blobs</a>.</p> <p><a href="https://learn.microsoft.com/en-us/azure/aks/concepts-storage">https://learn.microsoft.com/en-us/azure/aks/concepts-storage</a></p>

Claim 1	Accused Instrumentalities
	<h2>Persistent volumes</h2> <p>Volumes defined and created as part of the pod lifecycle only exist until you delete the pod. Pods often expect their storage to remain if a pod is rescheduled on a different host during a maintenance event, especially in StatefulSets. A <i>persistent volume</i> (PV) is a storage resource created and managed by the Kubernetes API that can exist beyond the lifetime of an individual pod.</p> <p>You can use the following Azure Storage services to provide the persistent volume:</p> <ul style="list-style-type: none"> <li>• Azure Disk</li> <li>• Azure Files</li> <li>• Azure Container Storage</li> </ul> <p>As noted in the <a href="#">Volumes</a> section, the choice of Azure Disks or Azure Files is often determined by the need for concurrent access to the data or the performance tier.</p> <pre> graph TD     AKS[Azure Kubernetes Service (AKS) cluster] --&gt; PV[Persistent Volume]     PV --&gt; ADisk[Azure managed disks (Standard or Premium storage)]     PV --&gt; AFiles[Azure Files (Standard storage)]     </pre> <p><a href="https://learn.microsoft.com/en-us/azure/aks/concepts-storage">https://learn.microsoft.com/en-us/azure/aks/concepts-storage</a></p>

Claim 1	Accused Instrumentalities
	<p><b>What kind of disks are supported for Windows?</b></p> <p>Azure Disks and Azure Files are the supported volume types, and are accessed as New Technology File System (NTFS) volumes in the Windows Server container.</p> <p><a href="https://learn.microsoft.com/en-us/azure/aks/windows-faq">https://learn.microsoft.com/en-us/azure/aks/windows-faq</a></p> <h2>About storage drivers</h2> <p>To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2>Storage drivers versus Docker volumes</h2> <p>Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p>Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the <a href="#">volumes section</a> to learn how to use volumes to persist data and improve performance.</p> <p><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="692 202 1136 262">Images and layers</h2> <p data-bbox="692 300 1881 376">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="713 437 1501 763"># syntax=docker/dockerfile:1  FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py</pre> <p data-bbox="692 829 1959 1139">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="692 1160 1284 1192"><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>

Claim 1	Accused Instrumentalities
	<p>Each layer is only a set of differences from the layer before it. Note that both <i>adding</i>, and <i>removing</i> files will result in a new layer. In the example above, the <code>\$HOME/.cache</code> directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the <a href="#">Best practices for writing Dockerfiles</a> and <a href="#">use multi-stage builds</a> sections to learn how to optimize your Dockerfiles for efficient images.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an <code>ubuntu:15.04</code> image.</p>  <p>Container (based on <code>ubuntu:15.04</code> image)</p> <p><a href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</a></p>

Claim 1	Accused Instrumentalities
	<h2>Volumes</h2> <p>Volumes are the preferred mechanism for persisting data generated by and used by Docker containers. While <a href="#">bind mounts</a> are dependent on the directory structure and OS of the host machine, volumes are completely managed by Docker. Volumes have several advantages over bind mounts:</p> <p><a href="https://kubernetes.io/docs/concepts/storage/volumes/">https://kubernetes.io/docs/concepts/storage/volumes/</a></p> <h2>Container environment</h2> <p>The Kubernetes Container environment provides several important resources to Containers:</p> <ul style="list-style-type: none"> <li>• A filesystem, which is a combination of an <a href="#">image</a> and one or more <a href="#">volumes</a>.</li> <li>• Information about the Container itself.</li> <li>• Information about other objects in the cluster.</li> </ul> <p><a href="https://kubernetes.io/docs/concepts/containers/container-environment/">https://kubernetes.io/docs/concepts/containers/container-environment/</a></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="699 213 925 279">Images</h2> <p data-bbox="699 306 1575 463">A container image represents binary data that encapsulates an application and all its software dependencies. Container images are executable software bundles that can run standalone and that make very well defined assumptions about their runtime environment.</p> <p data-bbox="699 496 1584 572">You typically create a container image of your application and push it to a registry before referring to it in a Pod.</p> <p data-bbox="677 597 1374 633"><a href="https://kubernetes.io/docs/concepts/containers/images/">https://kubernetes.io/docs/concepts/containers/images/</a></p> <h2 data-bbox="699 675 967 740">Volumes</h2> <p data-bbox="699 776 1584 1225">On-disk files in a container are ephemeral, which presents some problems for non-trivial applications when running in containers. One problem occurs when a container crashes or is stopped. Container state is not saved so all of the files that were created or modified during the lifetime of the container are lost. During a crash, kubelet restarts the container with a clean state. Another problem occurs when multiple containers are running in a Pod and need to share files. It can be challenging to setup and access a shared filesystem across all of the containers. The Kubernetes volume abstraction solves both of these problems. Familiarity with Pods is suggested.</p> <p data-bbox="677 1250 1364 1286"><a href="https://kubernetes.io/docs/concepts/storage/volumes/">https://kubernetes.io/docs/concepts/storage/volumes/</a></p>

Claim 1	Accused Instrumentalities
	<p><b>Open Container Initiative</b></p> <hr/> <p><b>Image Format Specification</b></p> <hr/> <p>This specification defines an OCI Image, consisting of an <a href="#">image manifest</a>, an <a href="#">image index</a> (optional), a set of <a href="#">filesystem layers</a>, and a <a href="#">configuration</a>.</p> <p>The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p> <p><a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</a></p>

Claim 1	Accused Instrumentalities
	<h2>Overview</h2> <p>At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more <a href="#">filesystem layer changeset</a> archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p> <pre> public class HelloWorld {     public static void main(String[] args) {         System.out.println("Hello, World");     } } </pre> <p><a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</a></p>

Claim 1	Accused Instrumentalities
	<h2>OCI Image Configuration</h2> <p>An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in <a href="#">Layers</a>.</p> <p>This section defines the <code>application/vnd.oci.image.config.v1+json</code> <a href="#">media type</a>.</p> <p><a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</a></p>

Claim 1	Accused Instrumentalities
	<p><b>Layer</b></p> <ul style="list-style-type: none"> <li>• Image filesystems are composed of <i>layers</i>.</li> <li>• Each layer represents a set of filesystem changes in a tar-based <a href="#">layer format</a>, recording files to be added, changed, or deleted relative to its parent layer.</li> <li>• Layers do not have configuration metadata such as environment variables or default arguments - these are properties of the image as a whole rather than any particular layer.</li> <li>• Using a layer-based or union filesystem such as AUFS, or by computing the diff from filesystem snapshots, the filesystem changeset can be used to present a series of image layers as if they were one cohesive filesystem.</li> </ul> <p><b>Image JSON</b></p> <ul style="list-style-type: none"> <li>• Each image has an associated JSON structure which describes some basic information about the image such as date created, author, as well as execution/runtime configuration like its entrypoint, default arguments, networking, and volumes.</li> <li>• The JSON structure also references a cryptographic hash of each layer used by the image, and provides history information for those layers.</li> <li>• This JSON is considered to be immutable, because changing it would change the computed <a href="#">ImageID</a>.</li> <li>• Changing it means creating a new derived image, instead of changing the existing image.</li> </ul> <p><a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</a></p>

Claim 1	Accused Instrumentalities
	<ul style="list-style-type: none"> <li>• <b>rootfs</b> <i>object</i>, REQUIRED           <p>The rootfs key references the layer content addresses used by the image. This makes the image config hash depend on the filesystem hash.</p> <ul style="list-style-type: none"> <li>◦ <b>type</b> <i>string</i>, REQUIRED               <p>MUST be set to <code>layers</code>. Implementations MUST generate an error if they encounter an unknown value while verifying or unpacking an image.</p> </li> <li>◦ <b>diff_ids</b> <i>array of strings</i>, REQUIRED               <p>An array of layer content hashes ( <code>DiffIDs</code> ), in order from first to last.</p> <p><a href="https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</a></p> </li> </ul> </li> </ul>